Claims 1-5, 7-15, 17, and 19-31 are now pending in the application. The Examiner is respectfully requested to reconsider and withdraw the rejections in view of the remarks contained herein.

REJECTION UNDER 35 U.S.C. § 103

Claims 1, 2, 5, 7-12, 17, 19 and 20 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Karp (U.S. Pat. No. 5,469,154) in view of Foster et al. (U.S. Pub. No. 2002/0181395).

Claims 3, 4, 13-15, and 21-23 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Karp (U.S. Pat. No. 5,469,154) in view of Foster et al. (U.S. Pub. No. 2002/0181395), and further in view of Brahmaroutu (U.S. Pub. No. 2003/0033427).

Claims 24, 25, and 28-31 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Karp (U.S. Pat. No. 5,469,154) in view of Dell et al. (U.S. Pub. No. 2002/0085578), and further in view of Foster et al. (U.S. Pub. No. 2002/0181395).

Claims 26 and 27 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Karp (U.S. Pat. No. 5,469,154) in view of Dell et al. (U.S. Pub. No. 2002/0085578), and further in view of Brahmaroutu (U.S. Pub. No. 2003/0033427).

These rejections are respectfully traversed.

Claim 1 recites that "wherein the forwarding instructions create a path between each of the plurality of sources and each of the plurality of destinations to make the CLOS network operate as a strictly non-interfering network." Applicant respectfully traverses the Examiner's assertion that Karp teaches the above limitations.

Karp at best appears to show a <u>circuit switched</u> network. No packet is switched. In other words, Karp is not concerned with the concept of, and cannot teach, "<u>non-interfering</u>" as defined in Applicant's specification where "the only queuing delays experienced by an admissible traffic pattern are attributable to the multiplexing of packets from slow links onto a faster link whose aggregate bandwidth at least equals the sum of the bandwidths of the smaller links."

Karp states, at col. 3, lns. 24-30,

> [u]sing a control algorithm for signal path revision, when necessary, this network allows each output port 9 to select any input port 6. The required non-interfering input-output paths will be set up regardless of the interconnection of other output-input ports. Each input port 6 can be selected by any number of output ports 9.

Even though Karp uses the term "non-interfering", it is only in the context of a circuit switched network and requires only the input-output paths do not interfere with each other (i.e., do not use the same circuit).

With respect to a packet switched network, Applicant's Specification states

> [0024] A non-interfering network (i.e. a network without interference) is a network for which the performance degradation for any admissible traffic pattern is guaranteed to conform to a pre-specified bound. This bound can be either deterministic or statistical. For example, a network can be deemed non-interfering if the worst-case end-to-end latency is guaranteed to be less than ten microseconds. This is an example of a deterministic bound. As another example, a network can be deemed non-interfering if 99% of packets experience network latencies of less than two microseconds. This is an example of a statistical bound. These are just examples and are not limiting of the invention. The appropriate choice for a pre-specified bound is application specific, and a network supporting multiple applications can impose different bounds on performance on each traffic type.

> [0025] <u>A strictly non-interfering network (SNIN) is a network for which the only queuing delays experienced by an admissible traffic pattern are attributable to the multiplexing of packets from slow links onto a faster link</u>

whose aggregate bandwidth at least equals the sum of the bandwidths of the smaller links. In a SNIN, competing traffic sources do not attempt to use the same network resources at the same time. The implementation of a SNIN requires that resources be dedicated through the network in support of an active communication session. In order to accomplish this, non-blocking networks can be used.

[0026] A network is non-blocking if it has adequate internal resources to carry out all possible admissible traffic patterns. There are different degrees of non-blocking performance based upon the sophistication of the control policy required to achieve non-blocking performance.

[0027] Most network switching applications allow the establishment of new connections and the tear down of old ones. It is possible that for a network with a non-blocking topology, a new connection can be blocked due to poor or unfortunate assignment of previously established connections. A strictly non-blocking network is a network for which any new admissible connection may be accepted independent of the state of preexisting connections, or the policy used to reroute preexisting connections, without changing the routes of the preexisting connections. A crossbar network is an example of a strictly non-blocking network. As another example, a rearrangably non-blocking network is a network that may be augmented by a mechanism to reroute preexisting connections such that it is possible to carry the preexisting connections and any new admissible connection.

Applicant's Substitute Specification of December 28, 2007, paras. [0024]-[0027] (emphasis added). The above paragraphs explain the difference between a strictly non-blocking network and a strictly non-interfering network. To emphasize, a strictly non-blocking network is a network for which any new admissible connection may be accepted independent of the state of preexisting connections, or the policy used to reroute preexisting connections, without changing the routes of the preexisting connections. A strictly non-interfering network (SNIN) is a network for which the only queuing delays experienced by an admissible traffic pattern are attributable to the multiplexing of packets from slow links onto a faster link whose aggregate bandwidth at least equals the sum of the bandwidths of the smaller links. In other words, a non-

blocking network <u>does not by itself operate as a strictly non-interfering network</u>. For example, a strictly non-blocking network may still allow queuing the competing packets to be switched at a switching node as long as "any new admissible connection may be accepted independent of the state of preexisting connections, or the policy used to reroute preexisting connections, without changing the routes of the preexisting connections."

The above point is further evidenced by Dell, the reference cited by the Examiner. Dell states

> [0006] In non-blocking switch fabrics, user data received from any particular source 102 can be routed through switch fabric 100 to any particular destination 110, <u>independent of whether any other user data is being routed from one or more other sources towards one or more other destinations through switch fabric 100</u>. <u>When packets of user data transmitted from the sources are received at an input port of input stage 104, the user data is buffered for eventual transmission through switching stage 106 to output stage 108, where the routed user data is again buffered for eventual transmission to the destination via an output port of the output stage associated with that destination</u>.

Dell, para. [0006] (emphasis added). In short, non-blocking switch fabrics may still allow competing packets to be buffered in a situation other than multiplexing of packets from slow links onto a faster link whose aggregate bandwidth at least equals the sum of the bandwidths of the smaller links.

Claim 1 further recites

> calculating a plurality of routing trees, each routing tree comprising the plurality of switches;
>
> calculating a plurality of Destination Location Identifiers (DLID) and a set of forwarding instructions for each of the plurality of first stage and second stage switches, wherein each of the plurality of DLIDs corresponds to one of the plurality of routing trees and one of the plurality of destinations; and

> populating a forwarding table of each of the plurality of first stage and second stage switches in the CLOS network with the plurality of DLIDs and the set of forwarding instructions[.]

The Examiner acknowledges that Karp fails to teach the above limitations, but asserts that one of ordinary skill in the art would be motivated to modify based on the teaching of Foster to arrive at the above limitations. Applicant respectfully traverses the Examiner's assertion.

Foster at best appears to show a Virtual Identifier (VI) Network Interface Controller (NIC) that receives an indication that a data communication to one or more remote nodes is to occur. The VI NIC will register the communication with a network manager for the network, and will receive an appropriate transmittal virtual identifier to be used for that communication from the network manager. Foster also states:

> [i]n some embodiments in which virtual identifiers are assigned to paths through a network, the assignment of paths to such virtual path identifiers is performed in a dynamic fashion after an indication is received that a data communication is to occur, such as by the network manager upon receipt of a data communication registration. The assigning of a virtual path identifier to a path can include the configuring of each of one or more intermediate routing devices (e.g., routers or switches) between the source and the destination, such as by the network manager, so that when one of the routing devices receives a data communication that includes the virtual identifier it will forward the communication in an appropriate manner either directly to the destination or instead to a next routing device along the path that is similarly configured.

Foster, para. [0048].

Again, Applicant respectfully submits that Karp at best appears to show a circuit switched network. The Examiner appears to suggest that one of ordinary skill in the art would be motivated to modify the circuit switched network based on techniques for a packet switched network shown in Foster. Applicant submits that because no packet is switched in the network of Karp, one of ordinary skill in the art would not know how to

apply the virtual identifiers, which accompany data packets, to the circuit switched network of Karp.

In view of the foregoing, Applicant submits that claims 1, 10, 17 and 24 define over the art cited by the Examiner. Likewise, claims 2-5, 7-9, and 21; claims 11-15 and 22; claims 19-20 and 23; as well as claims 25-31 which depend from respective claims 1, 10, 17, and 24 also define over the art cited by the Examiner for one or more of the reasons set forth above regarding claim 1.

CONCLUSION

It is believed that all of the stated grounds of rejection have been properly traversed, accommodated, or rendered moot. Applicant therefore respectfully requests that the Examiner reconsider and withdraw all presently outstanding rejections. It is believed that a full and complete response has been made to the outstanding Office Action and the present application is in condition for allowance. Thus, prompt and favorable consideration of this amendment is respectfully requested. If the Examiner believes that personal communication will expedite prosecution of this application, the Examiner is invited to telephone the undersigned at (248) 641-1600.

Respectfully submitted,

Dated: December 15, 2008

By: /Joseph M. Lafata/
Joseph M. Lafata, Reg. No. 37,166

HARNESS, DICKEY & PIERCE, P.L.C.
P.O. Box 828
Bloomfield Hills, Michigan 48303
(248) 641-1600

JML/PFD/evm